

Wavelet Multi-scale Edge Detection for Extraction of Geographic Features to Improve Vector Map Databases

Marlin L. Gendron

(*US Naval Research Laboratory*)

Dr. Juliette W. Ioup

(*University of New Orleans*)

Although numerous at smaller geographic scales, vector databases often do not exist at the more detailed, larger scales. A possible solution is the use of image processing techniques to detect edges in high-resolution satellite imagery. Features such as roads and airports are formed from the edges and matched up with similar features in existing low-resolution vector map databases. By replacing the old features with the new more accurate features, the resolution of the existing map database is improved. To accomplish this, a robust edge detection algorithm is needed that will perform well in noisy conditions. This paper studies and tests one such method, the Wavelet Multi-scale Edge Detector. The wavelet transform breaks down a signal into frequency bands at different levels. Noise present at lower scales smoothes out at higher levels. It is demonstrated that this property can be used to detect edges in noisy satellite imagery. Once edges are located, a new method will be proposed for storing these edges geographically so that features can be formed and paired with existing features in a vector map database.

KEY WORDS

1. Maps/Charts. 2. Database. 3. Wavelets.

1. INTRODUCTION. While ancient navigators struggled to locate features crudely noted on their hand-drawn maps, modern travellers enjoy the luxury of high-speed computers, global positioning systems (GPS), and aerospace technology. Early adventurers relied on the wind to carry their ships and on their eyes to spot land. Today, our eyes extend around the globe in the form of multi-million dollar satellites, and our current location is tracked by a symbol moving across a computer screen (Lohrenz et al., 1997).

From tourists using moving map GPS systems in some rental cars to on-board computers that guide smart weapons to their targets, the demand for accurate and reliable maps is increasing. There are many digital map databases available to the military, commercial sector, and the public. Off-the-shelf software offers worldwide coverage, and maps of almost every major city are available via the Internet. Without a doubt, maps are important, but maintaining and updating these map databases is costly and time consuming. In many cases, humans are still required to examine

Report Documentation Page			Form Approved OMB No. 0704-0188					
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>								
1. REPORT DATE JAN 2000	2. REPORT TYPE	3. DATES COVERED 00-00-2000 to 00-00-2000						
4. TITLE AND SUBTITLE Wavelet Multi-scale Edge Detection for Extraction of Geographic Features to Improve Vector Map Databases			5a. CONTRACT NUMBER					
			5b. GRANT NUMBER					
			5c. PROGRAM ELEMENT NUMBER					
6. AUTHOR(S)			5d. PROJECT NUMBER					
			5e. TASK NUMBER					
			5f. WORK UNIT NUMBER					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Code 7440.1, Stennis Space Center, MS, 39529			8. PERFORMING ORGANIZATION REPORT NUMBER					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)					
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)					
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited								
13. SUPPLEMENTARY NOTES Journal of Navigation 53:1, pp. 79-92. January								
14. ABSTRACT								
15. SUBJECT TERMS								
16. SECURITY CLASSIFICATION OF: <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33.33%; padding: 2px;">a. REPORT unclassified</td> <td style="width: 33.33%; padding: 2px;">b. ABSTRACT unclassified</td> <td style="width: 33.33%; padding: 2px;">c. THIS PAGE unclassified</td> </tr> </table>			a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified						

satellite imagery and digitise features such as roads and airports into points that form vector representations. These points are stored in new map databases or used to update existing ones (Edmonds, 1998).

This paper explores one possible method of improving the resolution of existing vector map databases by using semi-automated image processing techniques. Features from high-resolution satellite images are extracted via a Wavelet Multi-scale Edge Detector (WMED) and placed in a geographic bitmap (Gendron et al., 1997). From the bitmap, vectors are formed and identical features are located and updated in the map database.

2. TRANSFORMS. Many problems in science cannot be solved directly. Tools are required to alter these problems in such a way that proven mathematical techniques yield results. For example, certain operations performed on equations in the time domain are difficult and often impossible. A Fourier transformation, which transforms a function from the time domain to the frequency domain, makes many of these operations that are impossible in the time domain manageable by performing them in the frequency domain. After operations are completed in the frequency domain, an inverse transform provides the means to return to the time domain (Bracewell, 1986).

The Fourier transform operates on a signal in the time domain and generates coefficients in the frequency domain. These coefficients describe the composition of the signal in terms of infinite-length sine and cosine functions. An alternative to using such long waves is to use short waves called wavelets. Portions of the input signal that most resemble the wavelet produce the highest coefficients in the output (Strang and Nguyen, 1996). Like the Fourier transform, the wavelet transform and its inverse are defined for both continuous and discrete signals. The discrete wavelet transform (DWT) lies at the heart of the edge detector used on the satellite images studied in this paper. The elements in the input series are pixels whose magnitudes are between 0 and 255. The position of the element does not denote time, but rather indicates its location in the image.

The wavelet transform is explained more fully later in this paper, but the time-scale (spatial-scale) property that wavelets exhibit is worth mentioning now. By using a pair of short wavelet filters and applying them to an input signal, *level one coefficients* are produced whose position in time relates directly back to the input. The wavelet filters are then rescaled, shifted, and applied to the output to produce *level two coefficients*. These coefficients are lower in resolution, but maintain their position relative to the original time series (Misiti and Misiti, 1996).

The time-scale (spatial-scale) properties of the DWT can be used to extract true features from a noisy image, because random white noise present at high resolutions tends to smooth out at lower resolutions. Peaks detected in lower level coefficients, which are not present in upper levels, are discarded as noise. Peaks spatially close at all or most levels are retained as true edges. This is the basic concept behind the WMED (Hajj et al., 1996a).

3. DIGITAL MAPS. A map can be stored on a computer in digital form simply as a picture; for example, a digital raster map image file. Most map files originate from paper charts scanned into the computer by a scanner. When pixels in the image correspond to points on a surface, such as the earth, the image is geo-rectified or geo-

registered. Several map files stored together in a structured way form a map database. Digital raster map databases exist at many different resolutions, or map scales, and define how individual map files can be pieced together to form complete maps (Lohrenz, 1991).

Frequently, when these maps are viewed, the current position of the cursor in latitude and longitude is displayed on the screen. Cursor readout implies that the computer knows the location of every pixel in the image. This is true and accomplished by geo-rectifying the image. This geo-rectifying process occurs by taking control points, at known latitude and longitude points on the earth, to generate coefficients. These coefficients are used to compute new x and y positions in the image. The number of control points and their precision determines the accuracy of the raster map. The geographic extent of the image and its size in pixel space determines the map scale of the image (Hager, et al., 1990; Landrum, 1989).

Raster maps make a lot of sense to the human brain. People are able to extract features – such as roads and highways – intuitively and determine their location relative to other features. Although the computer ‘knows’ the location of each pixel at a given resolution, the computer is unable to extract, identify, and determine the location of a specific feature (such as a road) easily.

An alternative to storing a raster picture is to store a vector database. A vector map database stores individual features, such as airports, roads and highways as points. A road feature is stored in the database as a set of points, including a beginning and ending position and intermediate points along its path where turns occur. The road’s position, magnitude and direction are reconstructed from these points. A digital vector map database contains many feature types and is considered accurate to a given map scale (Edmonds, 1998). Today, most features in a vector database are extracted from raster maps or paper charts, which is a costly procedure. When more accurate information (e.g., higher resolution satellite imagery) becomes available, the vectoring process must start afresh to create a new vector database or correct an existing one. A common method of correcting an existing database is to update individual features or feature layers. This paper will explore a method of updating existing features by using a WMED.

4. WAVELETS. Fourier analysis transforms a signal based in time to one based in frequency. Although Fourier analysis is an extremely useful tool, it has one serious disadvantage. Time information disappears when a signal is transformed to the frequency domain. When analysing the frequency component, it is impossible to know when an event occurred in time (Bracewell, 1986). The advent of the Short-Time (or Finite) Fourier Transform helped to solve the problem by windowing the signal and mapping it into a function of time and frequency. Although a step in the right direction, the chosen window size remains the same for all frequencies. This is a serious drawback. A new windowing technique solved the problem by allowing smaller windows for higher frequencies and larger windows for lower frequencies. This new technique, called wavelet analysis, uses the wavelet transform to break down signals into frequency bands at different scales or levels (Misiti and Misiti, 1996).

A wavelet transform contains a filter bank that separates a signal into high and low frequency bands by passing the signal through a highpass and lowpass filter. The highpass filter is a moving difference filter capable of detecting sharp changes in the data, and the lowpass filter is a moving average filter that smoothes the data. The




Figure 1. Haar Wavelet rescaled and shifted.




Figure 2. One Level DWT.

output of the lowpass and highpass filters are coefficients referred to as the ‘approximation’ and ‘detail’ coefficients, respectively (Strang and Nguyen, 1996).

A wavelet is a short wave with an average value of zero. By scaling the wavelet with a scale factor and shifting its position, approximation and detail coefficients can be generated at different levels. For the DWT, the position and scale factor are based on powers of two. This is called a dyadic approach (Strang and Nguyen, 1996). The scale factor, derived from the scaling function, is associated with the lowpass filter. The actual wavelet corresponds to the highpass filter. The formula that describes this relationship and produces the scaling function and the wavelet equation is known as the dilation equation. The dilation equation is a two-scale difference function that governs the wavelet’s ability to scale and shift. For example, dilating and shifting a simple boxcar function creates the Haar wavelet (Morgan, 1997). Both the Haar wavelet and its rescaled version, shown in Figure 1, exhibit an extraordinary ability to find edges in noisy signals (Hadj et al., 1996a).

Figure 2 shows a one-level DWT. The lowpass filter produces approximation coefficients by convolving the input signal with lowpass filter coefficients derived from the dilation equation. Reversing the lowpass filter and alternating the signs of the filter coefficients produces the highpass filter. Convolving this filter with the input signal results in the detail coefficients (Misiti and Misiti, 1996).

The number of approximation coefficients plus the number of detail coefficients is twice that of the input. To avoid doubling the number of terms after each level, the coefficients are down-sampled by a factor of two at the output of each filter to remove redundant information. Down-sampling by a factor of two means every other term is removed. In any given filter output, the lost terms normally cannot be recovered,




Figure 3. Two Level DWT Tree.




Figure 4. Five Level 1-D WMED.

but remaining components from both filters can always be used to recover fully the original signal (Strang and Nguyen, 1996).

More levels are obtained by iteration. Successive approximations are decomposed so that one signal is broken down into many lower resolution components or levels. Figure 3 shows this wavelet decomposition tree.

By up-sampling the coefficients and reversing the filtering process, the signal is reconstructed. Up-sampling is accomplished by adding zeros between every term, thus taking the place of the coefficient removed by down-sampling. Although half of the coefficients are replaced by zeros at each level, perfect reconstruction (PR) occurs if the coefficients are not manipulated in any other way, and the same wavelet is used. PR means the original output signal is returned unaltered (Strang and Nguyen, 1996). Reconstruction, via the Inverse Discrete Wavelet Transform (IDWT), is not needed for the WMED studied in this paper. Only the detail coefficients that contain the high frequency information are used to find the edges. The magnitude of the detail coefficients relates directly to how closely that part of the signal resembles the wavelet at the given level. At any given level, the number of large coefficients (those with high energy) is small (Strang and Nguyen, 1996).

Groups of large magnitude detail coefficients, called local maxima, possibly define edges (Hwang and Chang, 1996). Small coefficients are likely to represent noise and should be removed using wavelet shrinkage or thresholding. There are two types of thresholding: hard and soft. Hard thresholding is accomplished by choosing a threshold value and setting all coefficients that fall below the threshold to zero (Laine and Zong, 1996). Alternatively, in soft thresholding, coefficients below the threshold are set to zero, but the larger coefficients are manipulated for good compression performance or noise reduction (Laine and Zong, 1996). Although soft thresholding

is preferable for most applications, hard thresholding is used for the WMED implemented in this paper because the final output of the detector is binary (Gendron et al., 1997).

For this paper, coefficients below the threshold are set to 0, and coefficients above are changed to a value of 1. If the new output value is 1, an edge exists, otherwise it does not. The edge information encoded in this manner can easily be stored in a bitmap. Later in this paper, a novel approach is introduced to equate an edge bit stored in a bitmap to a geographic location on the earth (Gendron et al., 1997).

5. ONE-DIMENSIONAL SIGNALS. The first step in designing a WMED is to determine the filter length of the level one highpass filter (Figure 4). If the filter is too short in length, the resolution will be too high, and some edges will be missed. The Haar wavelet, used in the DWT for this paper, has only two coefficients. By lengthening the wavelet, or adding more coefficients, the filter overlap with the data is increased, and the resolution reduced (Hajj et al., 1996a).

To determine the optimal length of the Haar highpass filter for this research, a simple one-dimensional (1-D) block signal, containing 512 samples, was run through a one-dimensional DWT using a Haar wavelet of length 4. At level one, the high frequency coefficients should reflect all the edges for the noiseless case via local maxima. An examination of the output detail coefficients revealed that some edges were missing. The resolution was too high. The length of the filter was modified using upsampling, which increases the filter length and doubles the frequency response. (Strickland, 1997). Adding zeros between each filter coefficient increases the filter length from two to four terms. The original signal was run through the DWT with the new length. Now all edges were reflected in the level one output.

With the length problem addressed, the block signal was run through a five level WMED. Figure 5 shows the block function and detail coefficients produced by the




Figure 5. Five Level DWT of Noiseless Signal.




Figure 6. Local Maxima of Detail Coefficients.




Figure 7. Output of 1-D WMED.

DWT at each level. At the highest resolution, or lowest level, d_1 , the peaks are sharp and well defined, but become broader and wider at higher levels. However, even at the lowest resolution, d_5 , the edges are easily detected by local maxima.

Next, local maxima are found at each level by hard thresholding to a value of 0 or 1. Figure 6 shows a plot of the binary sequence of indicators at each level. Notice most indicators match up with the step edges at all levels.

The rule for tracing edges across levels is to connect an indicator at the lowest level to the spatially closest indicator at the higher levels. If indicators are found at most or all levels, it is considered an edge. Indicators at the lowest level that do not have spatially close matches at upper levels are ignored (Hajj et al., 1996a). The smallest step in the input signal, the seventh edge from the leftmost point at $t = 0$, is detected at level one, but lost in subsequent levels. It is possible to detect this edge by lowering the threshold value. This is acceptable in this case, but the relaxation of the threshold can have profound implications in noisy conditions.

Figure 7 shows the final output of the one-dimensional WMED after indicators are matched up. Notice that the seventh edge is not detected.




Figure 8. Five Level DWT of Signal with Noise.




Figure 9. WMED Indicators for Signal with Noise.

6. NOISE. To determine how well the WMED would perform in noisy conditions, Gaussian noise is added to the same one-dimensional (1-D) block function. Figure 8 shows the detail coefficients from the output of the DWT at five levels. Notice the positions of the edges are far more difficult to discern.




Figure 10. Output from 1-D WMED.




Figure 11. 2-D WMED.

Applying the DWT to the noisy signal, detail coefficients at level one (d_1) reflect edges and noise. It is not until higher levels, around d_3 , that the noise subsides and true edges emerge. Figure 9 shows the local maxima at five levels. The indicators are found by thresholding the maxima to values of 1 or 0.

Next, indicators, or possible edges, at d_1 are matched up with spatially close indicators at higher levels (d_2-d_4). There are two approaches to the matching process, and both are *ad-hoc*. In one method, features are reported if indicators are found at most levels. Another approach detects a feature when indicators are found at all the levels (Hajj et al., 1996a). Because a higher threshold value has been used to reduce noise, only the first four levels are examined. Figure 9 adds credibility to this decision, since the lowest resolution level (d_5) shows some edge dropout. The resulting edges are shown in Figure 10. The results are identical to the noiseless case, including the absent seventh step edge.

7. TWO-DIMENSIONAL IMAGES. The WMED can be extended to two dimensions (2-D) by scanning individual rows and columns of an image. Each scanline passes through the one-dimensional WMED to form a binary map. Row scanlines detect vertical edges, and column scanlines detect horizontal edges. By combining the two sets of edges into a two-dimensional binary map, all possible edges emerge including diagonal ones. Figure 11 shows the 2-D WMED (Hajj et al., 1996b).

If the two-dimensional input image happens to be a geo-rectified raster map image, edges in the output equate to geographical locations in the image. By knowing the

start coordinates of the scanline and all pixels that compose it, the geographic location of the WMED output can be calculated. Because of down-sampling differences, map scale issues, and coordinate system considerations, the process is non-trivial (Hager et al., 1990). To reduce the complexity, a geographic bitmap provides the mechanism to geo-rectify the edge information. In the bitmap, every bit represents a unique location on the earth. Set bits denote edges and equate to specific coordinates on the earth. Although the bitmap is defined for the entire world at a given map scale, memory is only allocated for bits that exist. This makes the bitmap scheme fast and compact (Gendron et al., 1997).

7.1. Digital Chart of the World Vector Database. The Digital Chart of the World, DCW, is a vector-based digital database produced as a navigational aid by the National Imagery and Mapping Agency (NIMA). The vector database contains worldwide coverage of a variety of map features including roads, utility lines, and railroads to name a few. The DCW database used in this research is only accurate at the 1:1,000,000 map scale. This is a small map scale and not intended for precise navigation (NIMA, 1992).

An intersection of a highway overpass and a road was extracted from the DCW vector database and displayed on the computer screen at the 1:1,000,000 map scale. The displayed image was saved into a raster image file called a Geographic Tagged Image Format (GTIF). GTIF is an extension of the popular Tagged Image Format (TIF) file. Unlike TIF, GTIF supports storing of map coordinates, datums, and projections. The raster image of the DCW vector feature is shown in Figure 12.




Figure 12. Digital Chart of the World.

The GTIF image is saved in 8-bit greyscale, which supports shades of grey ranging from 0 (black) to 255 (white). In this case, the features in the image are black, and the background is pure white, so the only colour indices required are 0 and 255. Replacing the 0 values with 1's and the 255 values with 0's, produces a binary feature map. The binary map and the geographical information are placed into a geographic bitmap and initialised at the 1:1,000,000 map scale. In other words, every pixel has a geographic location, and, in a sense, the vectors are frozen at this map scale (Gendron et al., 1997).

7.2. High-Resolution Satellite Images. A high-resolution 1:50,000 satellite image of the same intersection is also stored in a GTIF file (Figure 13). At this resolution,



Figure 13. 1:50K Geo-registered satellite image.



Figure 14. 1:1M DCW overlaying 1:50K satellite image.

individual halves of the divided roads can be seen, and the geographical location of the features is more accurate than the 1:1,000,000 vector data of the same features.

Since the vector features are represented in the bitmap, and the satellite raster image is geo-registered via the GTIF, the two can be merged together geographically. The resulting GTIF is shown in Figure 14.

It is not surprising that the features do not match up spatially, because the satellite image is a higher resolution and therefore more accurate. Edges from the original satellite image, Figure 13, can be extracted to improve the DCW vector data. These edges are grouped into features to form new vectors. To find the edges, the original satellite GTIF, Figure 13, is separated into row and column scanlines, and then run through the 2-D WMED. The output binary feature map is placed in a 1:50,000 geographic bitmap (Figure 15).



Figure 15. Geographic Bitmap of 2-D WMED Output.

With the edge information detected and stored in the geographic bitmap, the set bits can be grouped together to form vector features. This is not a trivial task, but there are several methods that can be implemented, including statistical analysis and artificial intelligence (AI). Either of these methods is beyond the scope of this paper.

In this experiment, the DCW vector features are extracted first from the vector database. In a more realistic implementation, new vector features obtained from the satellite imagery would be used to search the DCW database geographically. Once likely vector candidates are found, an object-matching algorithm would be implemented to identify and replace the old vector features with the new (Strickland and Hahn, 1996). Again, such an algorithm is not studied here, but Figure 16 shows a block diagram of the entire process.




Figure 16. Update Process.

8. CONCLUSION. The method for extraction of geographic features presented relies on a robust edge detector. The detector must not only find edges, but find them well in noisy conditions. A relatively new tool in image processing, the wavelet transform, lends its multi-resolution capabilities to the problem. Unlike the Fourier transform, use of wavelets allows analysis of signals in the frequency domain while retaining time or spatial information. This time-scale (spatial-scale) property of wavelets is utilized in the Wavelet Multiscale Edge Detector (WMED) to detect edges in noisy conditions.

The WMED was developed in three steps and tested on three data sets. The first step is to determine which wavelet best detects edges when incorporated into the WMED scheme. The wavelet must have discrete filters for computer digital processing, and the choice of scales is dyadic. The Daubechies 4 and 6 wavelets were tested (results not shown in this paper), but the Haar wavelet was chosen for its edge detection abilities in the presence of noise. This decision was supported by previous wavelet multiscale edge detection research (Hajj, Nguyen and Chin, 1996a).

The second step was to determine the proper length of the level one highpass filter. In a noiseless signal, the WMED must detect all edges. To ensure this, indicators of the feature must be located at all levels, especially in the first level. The rule for tracing edges across levels is to connect an indicator at the lowest level to the indicator at the same location at the higher levels. If indicators are found at all levels, it is considered an edge. Therefore, all edges must be detected in the level one detail coefficients or edges are lost.

To determine the correct filter length, the first data set used was a simple noiseless one-dimensional block signal containing 512 terms. Tests showed that the level one Haar highpass filter was too short and some edges were missed in the detail coefficients. The length of the highpass filter was increased from two coefficients to four by upsampling to correct the length problem. The new filter was implemented, and indicators occurred in all levels. All edges were present in the final output.

The last step involved deciding what type of thresholding to use and the best threshold value. To determine this, random white noise was added to the 512 block signal. Small coefficients at each level likely represent noise. To remove these coefficients, hard thresholding was used. Hard thresholding was chosen because the final output is a bitmap with values of 1 or 0. The best results for the one-dimensional signals are produced by discarding, or setting to zero, all coefficients below 60% of the maximum absolute value, and setting the remaining coefficients to 1.

A two-dimensional high-resolution 1:50,000 satellite image was run through the WMED using the Haar wavelet with the previously determined optimum filter length.

A hard threshold value of 60% of the maximum absolute value again produced the best results, finding the edges of the image reasonably well. The edges were stored in a two-dimensional geographic bitmap, which provided the mechanism to geo-rectify the edge information. Finally, the process of grouping these set bits together, forming vector representations, and improving the low-resolution database was discussed.

ACKNOWLEDGMENTS

This paper presents results of research conducted for the author's Master of Science thesis at the University of New Orleans Physics Department. We wish to acknowledge support from the NRL science and technology program and express our appreciation for the support of Dr. Eric Hartwig, NRL Code 7000 Superintendent.

REFERENCES

- Bracewell, R. N. (1986). *The Fourier Transform and Its Applications*. McGraw-Hill, second edition revised.
- Burrus, C., Gopinath, R. and Guo, H. (1998). *Introduction to Wavelets and Wavelet Transforms*. Prentice Hall.
- Canny, J. (1986). A computational approach to edge detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, **8**, 679–698.
- National Imagery and Mapping Agency (1992). *Digital Chart of the World (DCW)* product specification.
- Dixon, T. and Sibul, L. (1996). Detection and classification of acoustic transients using wavelet transforms. *SPIE*, 2825, pp. 342–351.
- Edmonds, D. (1998). Feedback from users of electronic chart technology. This *Journal* **52**, 144–148.
- Gendron, M. L., Wischow, P. B., Trenchard, M. E., Lohrenz, M. C., Riedlinger, L. M. and Mehaffey, M. J. (1997). Moving map composer. *Naval Research Laboratory*, Patent Pending, Case No. 76,358, 1–29.
- Hager, J., Fry, L., Jacks, S. and Hill D. (1990). Datums, ellipsoids, grids, and grid reference systems. *Defense Mapping Agency*, TM8358.1.
- Hajj, H., Nguyen, T. and Chin R. (1996a). On multi-scale feature detection using filter banks. *IEEE Trans. on Sig. Proc.*, **1**, 70–74.
- Hajj, H., Nguyen, T. and Chin R. (1996b). A 2-D multi-rate Bayesian framework for multi-scale feature detection. *Proceedings of the SPIE – The International Society for Optical Engineering*, 2825, pp. 330–341.
- Hwang, W. and Chang F. (1996). Character extraction from documents using wavelet maxima. *SPIE*, 2825, pp. 1003–1015.
- Laine, A. and Zong, X. (1996). A multi-scale sub-octave wavelet transform for de-noising and enhancement. *SPIE*, 2825, pp. 238–249.
- Landrum, J. (1989). Illustration of the relationships among database resolution, map scale, and display technique. *Naval Research Laboratory*. NORDA Technical Note 439, pp. 1–7.
- Lohrenz, M. (1991). The Navy standard compressed aeronautical chart (CAC) database. *Naval Research Laboratory*, SP 024:351:91, pp. 1–93.
- Lohrenz, M., Trenchard, M., Myrick, S., Perniciaro, R. and Gendron M. (1997). Digital map requirements study in support of advanced cockpit moving map displays. *Naval Research Laboratory*, NRL/FR/7441–96–9652, pp. 1–72.
- Mallat, S. and Zhong, S. (1992). *Wavelet Transform Maxima and Multi-scale Edges*. Jones and Bartlett Publishers, pp. 67–103.
- Misiti, M. and Misiti, Y. (1996). *Wavelet Toolbox*. The Math Works Inc.
- Morgan, D. (1997). Haar Wavelet. *Embedded Systems*, **10**, 91–96.
- Strang, G. and Nguyen, T. (1996). *Wavelets and Filter Banks*. Wellesley-Cambridge Press.
- Strickland, R. (1997). Wavelet transform methods for object detection and recovery. *IEEE Trans. on Image Processing*, **6**, 724–735.
- Strickland, R. and Hahn, H. (1996). Wavelet methods for extracting objects from complex backgrounds. *IEEE Southwest Symposium on Image Analysis and Interpretation* pp. 7–12.